

Policy Optimization in Robust Control: Weak Convexity and Subgradient Methods

Yuto Watanabe, Feng-Yi Liao, and Yang Zheng

Abstract—Robust control seeks stabilizing policies that perform reliably under adversarial disturbances, with $\mathcal{H}_\infty$ control as a classical formulation. It is known that policy optimization of robust $\mathcal{H}_\infty$ control naturally lead to nonsmooth and nonconvex problems. This paper builds on recent advances in nonsmooth optimization to analyze discrete-time static output-feedback $\mathcal{H}_\infty$ control. We show that the $\mathcal{H}_\infty$ cost is *weakly convex* over any convex subset of a sublevel set. This structural property allows us to establish the first non-asymptotic deterministic convergence rate for the subgradient method under suitable assumptions. In addition, we prove a weak Polyak–Łojasiewicz (PL) inequality in the state-feedback case, implying that all stationary points are globally optimal. We finally present a few numerical examples to validate the theoretical results.

I. INTRODUCTION

Ensuring robustness against unknown disturbances is crucial in control systems. In this context, $\mathcal{H}_\infty$ control has served as a cornerstone in robust control [1], which aims to design a stabilizing policy that minimizes the $\mathcal{H}_\infty$ norm of the closed-loop system. It has long been recognized that policy optimization for robust control naturally leads to nonconvex and nonsmooth optimization problems [2]–[8]. While Riccati equation-based and LMI-based methods have been developed to circumvent these difficulties [9], [10], they are typically limited to unstructured and model-based controller design. To address this, some recent works have proposed local search algorithms that directly optimize the original nonconvex nonsmooth problem [2]–[4], [11]–[13]. In particular, the authors in [2], [13] have implemented two official MATLAB solvers, `hinfstruct` and `systune`, for structured $\mathcal{H}_\infty$ control, with successful real-world applications (e.g., a space probe by the ESA [14]).

However, owing to the inherent nonconvexity and nonsmoothness, the theoretical understanding remains limited, particularly regarding the non-asymptotic convergence of local search algorithms. The works [2], [12], [13] proposed *deterministic* local search algorithms with some structure exploitation, but only admit *at most* asymptotic convergence. As a *probabilistic* approach, [11] presented an open-source package `HIFOO` based on a randomized gradient sampling strategy [15], which was also limited to asymptotic guarantees. Recently, [3], [4] presented algorithms with a non-asymptotic convergence rate in a probabilistic sense using recent randomization-based techniques in [16], [17]. However,

such probabilistic guarantees are subject to a risk of failure, which may lead to concerns in safety-critical systems. To our knowledge, deterministic non-asymptotic convergence rates are still largely open for $\mathcal{H}_\infty$ optimization.

This work addresses the aforementioned issue of convergence rates via *weak convexity*, and establishes a non-asymptotic convergence rate for the simple subgradient method. Recent studies [18]–[21] demonstrate that weakly convex functions are arguably one of the broadest classes in nonsmooth nonconvex optimization that admits deterministic non-asymptotic rates, which cover various practical functions (e.g., convex, smooth, the composite of a convex function and a smooth map). Building on the recent advances, we first establish weak convexity of the $\mathcal{H}_\infty$ cost for discrete-time systems with static output-feedback. While previous works [13], [22] pointed out a relevant property, called *lower- C^2* [23], we demonstrate that the $\mathcal{H}_\infty$ cost is more structured than lower- C^2 . With this property, we present a non-asymptotic convergence rate for the subgradient method under the assumption that the cost remains finite. Further, we present a *weak Polyak–Łojasiewicz* (PL) inequality [24] for the full state-feedback case, ensuring the global optimality of stationary points despite the nonconvexity.

Our major technical contributions are threefold:

- 1) **Weak convexity.** Since the stability constraint on feedback gains is inherently nonconvex, the standard definition of weak convexity is not applicable. Thus, we instead show *m*-weak convexity with some $m > 0$ over any convex subset of a sublevel set. Notably, the constant m is uniform over the sublevel set, which allows for a straightforward extension of the standard weakly convex case [18];
- 2) **A weak Polyak–Łojasiewicz (PL) inequality.** In the full state measurement case, we present a weak PL inequality [24], guaranteeing the global optimality of stationary points. Unlike the differentiability-based result in [24], our proof applies to the nonsmooth $\mathcal{H}_\infty$ cost. The argument leverages an analogous idea to the extended convex lifting [7] to analyze the nonsmooth $\mathcal{H}_\infty$ landscape beyond standard stationarity;
- 3) **The subgradient method.** Building on weak convexity, we extend the convergence analysis of [18] to the constrained nonconvex $\mathcal{H}_\infty$ optimization setting. In particular, we show that an ϵ -inexact stationary point can be obtained after $T = \mathcal{O}(1/\epsilon^4)$ iterations, provided the cost remains bounded along the iterates. To our knowledge, this is the first deterministic complexity guarantee for subgradient methods in robust control.

This work is supported by NSF ECCS-2154650, NSF CMMI 2320697, and NSF CAREER 2340713.

Y. Watanabe, F.-Y. Liao, and Y. Zheng are with the Department of Electrical and Computer Engineering, University of California San Diego; {y1watanabe, fliao, zhengy}@ucsd.edu.

The remainder of this paper is organized as follows. [Section II](#) formulates the $\mathcal{H}_\infty$ optimization problem and provides the problem statement. Then, [Section III](#) presents a review of landscape properties and our main results, such as the weak convexity and weak PL condition. In [Section IV](#), we discuss the subgradient method. [Section V](#) provides numerical results. Finally, [Section VI](#) concludes this paper.

II. PRELIMINARIES AND PROBLEM STATEMENT

A. Robust $\mathcal{H}_\infty$ optimization with static output-feedback

Consider a discrete-time linear time-invariant system¹

$$x_{t+1} = Ax_t + Bu_t + B_w w_t, \quad y_t = Cx_t, \quad (1)$$

where $x_t \in \mathbb{R}^{n_x}$ is the system state, $u_t \in \mathbb{R}^{n_u}$ is the control input, $y_t \in \mathbb{R}^{n_y}$ is the output measurement, and $w_t \in \mathbb{R}^{n_w}$ is the disturbance, respectively. The system matrices are $A \in \mathbb{R}^{n_x \times n_x}$, $B \in \mathbb{R}^{n_x \times n_u}$, $B_w \in \mathbb{R}^{n_x \times n_w}$, and $C \in \mathbb{R}^{n_y \times n_x}$. The system's initial condition is fixed as $x_0 = 0$.

We consider $\mathbf{w} = \{w_0, w_1, \dots\}$ as an adversarial disturbance with bounded energy. Let ℓ_2^k be the set of square-summable (bounded energy) signals of dimension k , i.e.,

$$\ell_2^k = \left\{ \mathbf{w} = \{w_0, w_1, \dots\} \mid \|\mathbf{w}\|_2 := \sqrt{\sum_{t=0}^{\infty} w_t^\top w_t} < \infty \right\}.$$

Robust $\mathcal{H}_\infty$ control aims to select the control sequence $\mathbf{u} = \{u_0, u_1, \dots\}$ to minimize the quadratic performance $\sum_{t=0}^{\infty} (x_t^\top Q x_t + u_t^\top R u_t)$ against the worst disturbance of bounded energy $\|\mathbf{w}\|_2 \leq 1$. Without loss of generality, we consider the energy bound on $\mathbf{w}$ to be 1, and we can formulate robust $\mathcal{H}_\infty$ control with any $\mathbf{w} \in \ell_2^{n_w}$. We thus consider the following min-max robust control problem

$$\min_{\mathbf{u}} \max_{\|\mathbf{w}\|_2 \leq 1} \sum_{t=0}^{\infty} (x_t^\top Q x_t + u_t^\top R u_t). \quad (2)$$

Throughout this paper, we make the following assumption

Assumption 1. *The system matrices (A, B) are stabilizable, and B_w and C are row full-rank. The weight matrices Q, R are positive definite, and the controller has access only to the output sequence $\mathbf{y} := \{y_0, y_1, \dots\}$.*

Under [Assumption 1](#), if full state measurements are available (i.e., $C = I_{n_x}$ in (1)), it is well known that the optimal solution of (2) is a static state-feedback policy $u_t = Kx_t$, where $K \in \mathbb{R}^{n_u \times n_x}$ is a constant gain matrix [1]. The general output-feedback $\mathcal{H}_\infty$ problem, however, is substantially more challenging [25]. In this work, we restrict attention to *linear static output-feedback policies* of the form

$$u_t = Ky_t, \quad t = 0, 1, 2, \dots, \quad (3)$$

and study the following policy optimization problem:

$$\begin{aligned} \min_{K \in \mathbb{R}^{n_u \times n_y}} \max_{\|\mathbf{w}\|_2 \leq 1} \sum_{t=0}^{\infty} (x_t^\top Q x_t + u_t^\top R u_t) \\ \text{subject to } (1) \text{ and } (3), x_0 = 0. \end{aligned} \quad (4)$$

¹We focus on the class of discrete-time systems, since it is practical and might also be the simplest case in $\mathcal{H}_\infty$ optimization, due to the *coercivity*; see [4, Lem. 3.2] and [7, Example 2.3]. See also [Remark 1](#).

Given any feedback gain $K \in \mathbb{R}^{n_u \times n_y}$, the closed-loop dynamics from (1) and (3) becomes $x_t = (A + BKC)x_t + B_w w_t$, $x_0 = 0$. We denote the set of stabilizing gains as

$$\mathcal{K} := \{K \in \mathbb{R}^{n_u \times n_y} \mid \rho(A + BKC) < 1\}, \quad (5)$$

where $\rho(\cdot)$ denotes the spectral radius.

Assumption 2. *The set $\mathcal{K}$ is non-empty.*

If $K \in \mathcal{K}$, then the state trajectory $\mathbf{x} = \{x_0, x_1, \dots\}$ belongs to $\ell_2^{n_x}$ for any bounded disturbance $\mathbf{w} \in \ell_2^{n_w}$. Let a performance signal be $z_t = (Q + C^\top K^\top R K C)^{1/2} x_t$ and $\mathbf{z} = \{z_0, z_1, z_2, \dots\}$. Then, we have

$$\|\mathbf{z}\|_2^2 = \sum_{t=0}^{\infty} (x_t^\top Q x_t + u_t^\top R u_t) < \infty, \quad \forall \mathbf{x} \in \ell_2^{n_x}, \mathbf{w} \in \ell_2^{n_w}.$$

For any $K \in \mathcal{K}$, the closed-loop system defined by (1) and (3) can be viewed as a bounded linear operator mapping disturbances $\mathbf{w} \in \ell_2^{n_w}$ to performance signals $\mathbf{z} \in \ell_2^{n_x}$. We denote this linear operator as $\mathbb{T}_{zw}(K)$, and define the $\ell_2 \rightarrow \ell_2$ induced norm as

$$\|\mathbb{T}_{zw}(K)\| := \max_{\|\mathbf{w}\|_2 \neq 0} \|\mathbf{z}\|_2 / \|\mathbf{w}\|_2 \quad \forall K \in \mathcal{K}. \quad (6)$$

Since $\mathbb{T}_{zw}(K)$ is linear, it is not difficult to see that problem (4) is equivalent to $\min_{K \in \mathcal{K}} \|\mathbb{T}_{zw}(K)\|^2$. This linear operator (6) is expressed in the time domain, and it also admits a frequency-domain representation via transfer functions. Denote transfer function matrix $\mathbf{T}_{zw}(K, \omega)$ from $\mathbf{w}$ to $\mathbf{z}$ as

$$\mathbf{T}_{zw}(K, \omega) = \begin{bmatrix} Q^{1/2} \\ R^{1/2} K C \end{bmatrix} (e^{j\omega} I - (A + BKC))^{-1} B_w,$$

where $\omega \in [0, 2\pi]$ is a frequency variable. Its $\mathcal{H}_\infty$ norm is

$$\|\mathbf{T}_{zw}(K)\|_{\mathcal{H}_\infty} = \max_{\omega \in [0, 2\pi]} \sigma_{\max}(\mathbf{T}_{zw}(K, \omega)), \quad \forall K \in \mathcal{K} \quad (7)$$

where $\sigma_{\max}(\cdot)$ is the maximum singular value. The $\ell_2 \rightarrow \ell_2$ induced norm coincides with the $\mathcal{H}_\infty$ norm [1, Th. 4.4], i.e., $\|\mathbb{T}_{zw}(K)\| = \|\mathbf{T}_{zw}(K)\|_{\mathcal{H}_\infty}$, $\forall K \in \mathcal{K}$.

Problem (4) is thus often called $\mathcal{H}_\infty$ optimization with static feedback policies, which can be equivalently written as

$$J^* = \min_{K \in \mathcal{K}} J(K), \quad (8)$$

where $J(K) := \|\mathbf{T}_{zw}(K)\|_{\mathcal{H}_\infty}$ is defined in (7). We remark that problem (4) minimizes $J(K)^2$, whereas (8) minimizes $J(K)$. Although the objective values differ by a square, the two problems clearly share the same set of optimal solutions.

B. Problem statement

Robust $\mathcal{H}_\infty$ control is a cornerstone of control theory and has been studied extensively since the 1980s; see e.g., the classical textbook [1]. In the state-feedback setting, standard approaches solve (8) either through Riccati-based iterative methods [9] or via convex reformulations with Lyapunov variables [10]. The static output-feedback is substantially more challenging, and no method is known to compute its globally optimal solution. Both Lyapunov- and Riccati-based approaches rely on explicit knowledge of the system model.

More recently, several works have explored direct optimization over the policy space $\mathcal{K}$ using local search methods [2]–[5], [11]–[13]. These approaches are typically more scalable, and some of them are better-suited to model-free settings [26], [27]. However, direct policy optimization for $\mathcal{H}_\infty$ control in (8) is inherently a *nonsmooth* and *nonconvex* problem, even in the state-feedback setting. Because of these difficulties, the early work [2], [11]–[13] provided no convergence-rate guarantees, while more recent studies [3]–[5] established global optimality in the state-feedback case but relied on substantially more intricate algorithmic frameworks, often involving advanced randomized techniques [15]–[17].

The main challenge is that the $\mathcal{H}_\infty$ cost function is both nonconvex and nonsmooth, with the nonsmoothness stemming from the maximization of the largest singular value in (7). While it is known that this function is *subdifferentially regular* [2]–[5], such regularity is too broad to yield efficient algorithmic guarantees. At the same time, the function is more structured than a generic nonsmooth, nonconvex objective: it is naturally a *lower- C^2* function [23], as can be seen directly from its definition (7).

This observation raises two central questions in this paper:

- 1) What additional structure (beyond subdifferential regularity) can be identified in the nonsmooth and nonconvex landscape of $\mathcal{H}_\infty$ policy optimization?
- 2) Given this structure, what are the simplest algorithms that can provably handle the nonconvex, nonsmooth nature of the $\mathcal{H}_\infty$ policy optimization problem?

To address the first question, we will show that the $\mathcal{H}_\infty$ cost function admits a form of *weak convexity*, and in the state-feedback case, even satisfies a *weak Polyak–Łojasiewicz (PL)* condition. To address the second, we will demonstrate that classical first-order methods, such as the *simple subgradient method*, can achieve deterministic non-asymptotic convergence rates that match the best-known complexity for weakly convex functions [18]–[21], while avoiding randomized or overly complex procedures as in [3], [4].

III. WEAK CONVEXITY IN $\mathcal{H}_\infty$ OPTIMIZATION

We first summarize some basic landscape properties in (8), then establish the notion of weak convexity, and finally discuss fundamental properties of its stationary points.

A. Basic landscape properties

We here summarize the basic landscape properties of (8).

Lemma 1. *Suppose Assumptions 1 and 2 hold. Consider the $\mathcal{H}_\infty$ optimization problem with static linear policies in (8). Then the following statements hold:*

- 1) *The stabilizing set $\mathcal{K}$ is open and, in general, nonconvex. If full state information is available (i.e., $C = I_{n_x}$), the set $\mathcal{K}$ is always path-connected, but in the general output-feedback case, it can be disconnected.*
- 2) *The $\mathcal{H}_\infty$ cost function $J : \mathcal{K} \rightarrow \mathbb{R}_+$ is coercive:*

$$J(K) \rightarrow \infty \quad \text{as } K \rightarrow \partial\mathcal{K} \text{ or } \|K\|_F \rightarrow \infty.$$

- 3) *The function J is generally nonsmooth and nonconvex.*
- 4) *On any sublevel set $\mathcal{K}_\nu := \{K \in \mathcal{K} \mid J(K) \leq \nu\}$, J is β_ν -Lipschitz continuous for some constant $\beta_\nu > 0$.*

These landscape properties are known in the literature; see, e.g., [2]–[6], [28]. We do not reproduce the proofs here, but instead highlight several insights and illustrative examples.

First, it is straightforward to construct examples showing that $\mathcal{K}$ is nonconvex (as Examples 2 and 4 below). When $C = I_{n_x}$, path-connectivity of $\mathcal{K}$ follows from an equivalent convex reformulation, since convex sets are always path-connected. Second, the coercivity of J can be established by deriving a global quadratic lower bound; see [4, Lem. 3.2] for details. This property is crucial because it ensures that each sublevel set $\mathcal{K}_\nu := \{K \in \mathcal{K} \mid J(K) \leq \nu\}$ is compact. Third, J is naturally nonconvex since its domain $\mathcal{K}$ is already nonconvex. Its nonsmoothness arises from two sources in the definition (7): (i) the maximum singular value of a complex matrix, and (ii) the maximization over the frequency variable. Finally, while the local Lipschitz continuity of J is well known, identifying a uniform Lipschitz constant on a sublevel set $\mathcal{K}_\nu$ is more subtle. Previous works such as [3, Th. 5] and [4, Sec. 3] assume this β_ν -Lipschitz property directly. In fact, it can be rigorously established by explicitly characterizing the Clarke subdifferential of J .

We next present two examples to illustrate Lemma 1.

Example 1 (Nonsmoothness of function J). *Consider*

$$x_{k+1} = x_k + u_k + w_k, \quad y_k = x_k \quad (9)$$

with $u_k = kx_k$ and $Q = R = 1$. It is clear that $\mathcal{K} = \{k \mid |1 + k| < 1\} = (-2, 0)$. It is easy to see

$$\mathbf{T}_{zw}(k, \omega)^H \mathbf{T}_{zw}(k, \omega) = \frac{1 + k^2}{1 + (1 + k)^2 - 2(1 + k) \cos \omega},$$

and thus the $\mathcal{H}_\infty$ cost can be obtained explicitly as

$$J(k) = \begin{cases} \frac{\sqrt{1+k^2}}{|k|}, & k \in (0, -1) \\ \frac{\sqrt{1+k^2}}{k+2}, & k \in (-1, -2). \end{cases}$$

The function J is clearly nonsmooth at $k = -1$, as illustrated in Figure 1a. Moreover, $J(k) \rightarrow \infty$ as $k \rightarrow 0$ or $k \rightarrow -2$, which demonstrates the coercivity property.

Consider another state-feedback system

$$A = \begin{bmatrix} 0 & 2 \\ 4 & 0.2 \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad B_w = I_2 \quad (10)$$

with $Q = \text{diag}(1, 10^{-3})$ and $R = 1$. From the classical $\mathcal{H}_\infty$ theory, we can compute the optimal value of $J(K)$ to be $J^ \approx 8.327$. We plot the function J for $K = [k_1, k_2]$ in Figure 1b. We can see nonsmooth points on the landscape.*

Example 2 (Nonconvexity of domain $\mathcal{K}$ and function J). *Consider the following problem data:*

$$A = \begin{bmatrix} 1 - \alpha & -0.1 & -0.1 \\ 0.1 & 1 & 0 \\ 0 & 0.1 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0.1 \\ 0 \\ 0 \end{bmatrix},$$

$$B_w = I_3, \quad C = \begin{bmatrix} 0 & 1 & 1 \\ 1 & -1 & 1 \end{bmatrix}, \quad Q = 0.01I_3, \quad R = 0.01.$$




Fig. 1: Nonconvex and nonsmooth landscape in $\mathcal{H}_\infty$ optimization. (a)–(b) Nonsmoothness of the cost functions J in [Example 1](#); (c) Disconnectivity of $\mathcal{K}$ in [Example 2](#) with $\alpha = 0.13$; (d) Nonconvexity of J in [Example 2](#) with $\alpha = 0.13$.

For this system, the static output-feedback gain is of the form $K = [k_1, k_2]$. When $\alpha \approx 0.05 \sim 0.13$, the feasible region $\mathcal{K}$ is disconnected, which directly leads to the nonconvex landscape. We plot set $\mathcal{K}$ and the function J for $\alpha = 0.13$ in [Figure 1c](#) and [Figure 1d](#), respectively.

Remark 1 ($\mathcal{H}_\infty$ optimization in continuous-time systems). The continuous-time counterpart of (8) is likewise nonsmooth and nonconvex. The set of stabilizing gains remains nonconvex, is path-connected when $C = I_{n_x}$, and can be path-disconnected in general. However, unlike the discrete-time case, the coercivity of J does not necessarily hold in continuous time (see [7, Example 2.3]). As a result, its sublevel sets may fail to be compact, and it is unclear whether J is β_ν -Lipschitz continuous on such sets. $\square$

B. Lower- C^2 and weak convexity

We now establish a key technical result: although the $\mathcal{H}_\infty$ cost J is nonconvex, it is in fact *weakly convex*. A precise statement will be given later in [Theorem 1](#). Recall that a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is ρ -weakly convex if the function $x \mapsto f(x) + \frac{\rho}{2}\|x\|^2$ is convex [18], [20]. Weak convexity provides a broad framework that subsumes convex, smooth, and many composite nonconvex functions, and enables principled first-order algorithm design [18]–[21].

Another related notion is that of *lower- C^2* functions. Roughly speaking, a lower- C^2 function may be nonsmooth, but its nonsmoothness arises solely from taking a supremum of smooth C^2 components, making it significantly more structured than generic nonsmooth functions. Note that a function $f : \mathcal{D} \rightarrow \mathbb{R}$ is of class C^2 if it is twice differentiable on its open domain $\mathcal{D}$ and both ∇f and $\nabla^2 f$ are continuous. The formal definition of lower- C^2 functions is as follows.

Definition 1 (Lower- C^2 [23]). We say a function $f : \mathcal{D} \rightarrow \mathbb{R}$, defined on an open domain $\mathcal{D} \subset \mathbb{R}^n$, is lower- C^2 , if for each $\bar{x} \in \mathcal{D}$, there exist a neighborhood V of $\bar{x}$, a compact set $\mathcal{T}$, and a mapping $\varphi : \mathcal{T} \times V \rightarrow \mathbb{R}$ such that

$$f(x) = \max_{t \in \mathcal{T}} \varphi(t, x), \quad \forall x \in V \quad (11)$$

where, for each $t \in \mathcal{T}$, the function $\varphi(t, \cdot)$ is C^2 in x , and the mappings $\varphi(\cdot, \cdot)$, $\nabla_x \varphi(\cdot, \cdot)$, and $\nabla_x^2 \varphi(\cdot, \cdot)$ are jointly continuous on $(t, x) \in \mathcal{T} \times V$.

Note that the smooth functions $\varphi(t, \cdot)$ in (11) may depend on the point $\bar{x}$. Thus, a lower- C^2 function can be *locally*

represented as the maximum of a family of smooth C^2 functions. We will show that the $\mathcal{H}_\infty$ cost J is lower- C^2 .

Before doing so, let us recall some simple examples. First, any quadratic function is lower- C^2 , and so is the maximum of finitely many quadratics², e.g., $f(x) = \max_{t \in \{1, \dots, p\}} f_t(x)$. Second, the maximum eigenvalue function of symmetric matrices $\lambda_{\max}(\cdot)$ is also lower- C^2 , as

$$\lambda_{\max}(X) = \max_{u^T u = 1} u^T X u, \quad (12)$$

where each component $\varphi(u, X) = u^T X u$ is linear in X and hence C^2 . Finally, the maximum singular value of a real matrix $Y \in \mathbb{R}^{p \times m}$ is lower- C^2 , because

$$\sigma_{\max}(Y) = \max_{\substack{u \in \mathbb{C}^m, v \in \mathbb{C}^p \\ \|u\| = \|v\| = 1}} \Re(v^H Y u), \quad (13)$$

where $\Re(\cdot)$ represents the real part and, each component $\varphi(u, v, Y) = \Re(u^H Y v)$ is linear (hence C^2) in Y . These examples illustrate that many spectral functions naturally fall into the lower- C^2 class, which enables our analysis of the $\mathcal{H}_\infty$ cost below; recall its definition (7).

Formally, we have the following result.

Lemma 2. Suppose [Assumptions 1](#) and [2](#) hold. Consider the $\mathcal{H}_\infty$ cost function $J : \mathcal{K} \rightarrow \mathbb{R}_+$ defined in (8). Then the following properties hold:

- 1) J is a lower- C^2 function.
- 2) J is locally weakly convex: for every $K_0 \in \mathcal{K}$, there exist constants $m > 0$ and $r > 0$ such that

$$K \mapsto J(K) + \frac{m}{2} \|K\|_F^2 \quad (14)$$

is convex over $\mathbb{B}_r(K_0) := \{K \mid \|K - K_0\|_F \leq r\}$.

- 3) J is subdifferentially regular, i.e., its directional derivative coincides with its Clarke directional derivative at every point.

Proof. Similar to (13), we can rewrite the $\mathcal{H}_\infty$ cost as

$$J(K) = \max_{\omega \in [0, 2\pi]} \max_{\substack{u \in \mathbb{C}^{n_w}, v \in \mathbb{C}^{n_x + n_u} \\ \|u\| = \|v\| = 1}} \Re[v^H \mathbf{T}_{zw}(K, \omega) u].$$

Define $\varphi(\omega, u, v, K) := \Re[v^H \mathbf{T}_{zw}(K, \omega) u]$. It is easy to verify that φ is analytic in K over the open set $\mathcal{K}$, and jointly continuous in (ω, u, v, K) on $[0, 2\pi] \times \{u \in \mathbb{C}^{n_w} \mid$

²In this case, we regard $\mathcal{T}$ as a subset of a discrete topological space.

$\|u\| = 1\} \times \{v \in \mathbb{C}^{n_x+n_u} \mid \|v\| = 1\} \times \mathcal{K}$. Hence, we may represent J as

$$J(K) = \max_{(\omega, u, v) \in \Theta} \varphi(\omega, u, v, K), \quad (15)$$

where $\Theta := [0, 2\pi] \times \{u \in \mathbb{C}^{n_u} : \|u\| = 1\} \times \{v \in \mathbb{C}^{n_x+n_u} : \|v\| = 1\}$ is compact. By Definition 1, it follows that J is lower- C^2 . The local weak convexity 2) and subdifferential regularity 3) follow from [23, Th. 10.31 and Th. 10.33]. $\square$

The proof above is inspired by the continuous-time case [22, Lem. 9]. Note that the local weak convexity (14) is a natural property of lower C^2 functions [23, Th. 10.33]. However, the weak convexity constant m may depend on the point and its neighborhood, which may be hard to estimate.

It is worth noting that the smooth components in (12), (13) and (15) are *global*, in the sense that they do not vary with the choice of reference point K . This structure is considerably stronger than general lower- C^2 functions in Definition 1. Consequently, we may anticipate better-behaved properties for the $\mathcal{H}_\infty$ cost than those established in Lemma 2. This is indeed the case: we can establish a uniform weak convexity constant over any convex subset in a sublevel set of J .

Theorem 1 (Weak convexity). *Suppose Assumptions 1 and 2 hold. Consider the $\mathcal{H}_\infty$ cost function $J : \mathcal{K} \rightarrow \mathbb{R}_+$ defined in (8). Let $\nu > 0$ and define a sublevel set $\mathcal{K}_\nu := \{K \in \mathcal{K} \mid J(K) \leq \nu\} \neq \emptyset$. Then, for any nonempty convex subset $V \subset \mathcal{K}_\nu$, there exists a constant $m_\nu > 0$ such that the function*

$$K \mapsto J(K) + \frac{m_\nu}{2} \|K\|_F^2$$

is convex over V .

The proof is not difficult by recalling the fact that the composition of a convex function with an L -smooth function is weakly convex [18]. In our case, the maximum singular value function $\sigma_{\max}(\cdot)$ is convex on $\mathbb{R}^{p \times n}$. We need to handle the complex transfer function and the maximization over the frequency variable ω in (7). For this, we first use a well-known result from complex SDPs [29, Exercise 4.42] to equivalently eliminate complex numbers in $f(\cdot, \omega) := \sigma_{\max}(\mathbf{T}_{zw}(\cdot, \omega))$ for each fixed $\omega \in [0, 2\pi]$. Then, since $\mathcal{K}_\nu$ is compact thanks to the coercivity, over any convex subset $V \subset \mathcal{K}_\nu$, we can view $f(\cdot, \omega)$ as a composition of a convex function with a real-valued L_ν -smooth function with some $L_\nu > 0$, which is thus m_ν -weakly convex with some $m_\nu > 0$ independent of ω . Therefore, we have $J(K) + \frac{m_\nu}{2} \|K\|_F^2 = \max_{\omega \in [0, 2\pi]} \{f(K, \omega) + \frac{m_\nu}{2} \|K\|_F^2\}$ is convex on any convex subset of $\mathcal{K}_\nu$. A detailed proof is given in Section III-D.

Example 3 (Weak convexity). *Consider $f : (-1, \infty) \rightarrow \mathbb{R}_+$:*

$$f(x) = \sigma_{\max}(c(x)), \quad c(x) = \begin{bmatrix} 1 & x/2 \\ (1+x)^{-1} & -2 \cos x \end{bmatrix}.$$

This function is slightly different from the $\mathcal{H}_\infty$ cost (7), as it has no frequency variable. However, it captures the essential composition of $\sigma_{\max}$ with a smooth function. Indeed, on any bounded interval $V \subset (-1, \infty)$, the mapping $c(\cdot)$ is L_V -smooth for some $L_V > 0$, and thus $f(\cdot)$ is m_V -weakly convex




Fig. 2: Plots of f and $f + \frac{5}{4} \|\cdot - 1\|^2$ on $V = [0, 5]$ in Example 3.




Fig. 3: Spurious stationary points of J with the set $\mathcal{K}$ in the case of static output-feedback in Example 4 (the same system as Example 2 with $\alpha = 0.14$). It can be observed that $J(K)$ possesses not only a local minimum but also a saddle point, represented by the red dot.

on V with some $m_V > 0$. To illustrate, we plot $f(x)$ and $f(x) + \frac{5}{4} \|x - 1\|^2$ on $V = [0, 5]$ in Figure 2. Clearly, $f(x)$ is convexified by adding a quadratic perturbation $\frac{5}{4} \|x - 1\|^2$.

C. Stationary points and weak PL property

We next analyze the stationary points of $\mathcal{H}_\infty$ cost J . As is typical in nonconvex optimization, spurious stationary points such as local minima and saddle points may arise in the general output-feedback $\mathcal{H}_\infty$ setting.

Lemma 3. *Suppose Assumptions 1 and 2 hold. Consider the $\mathcal{H}_\infty$ cost function $J : \mathcal{K} \rightarrow \mathbb{R}_+$ defined in (8). When $C \neq I_{n_x}$, the function J may admit spurious stationary points, including local minima and saddle points.*

We verify this result using Example 4 below. This is the same instance as Example 2 with a different choice of α . The example highlights that such spurious stationary points naturally emerge from the nonconvexity and potential disconnectedness of $\mathcal{K}$.

Example 4 (Saddle points). *Consider the problem data in Example 2. For $\alpha = 0.13$, the set $\mathcal{K}$ consists of two disconnected components (see Figure 1c). At $\alpha = 0.14$, these components merge, and $\mathcal{K}$ becomes connected, as shown in Figure 3a. In this case, we observe a saddle point in J (see Figure 3b), as well as at least one local minimum.*

In the full-state measurement case ($C = I_{n_x}$), the $\mathcal{H}_\infty$ optimization exhibits a much more tractable landscape, with a connected feasible region and convex-like shape as Figures 1a and 1b. Indeed, we can establish a weak PL condition that ensures the global optimality of stationary points.

Theorem 2 (Weak PL condition). *Suppose Assumptions 1 and 2 hold. If we have full state measurement, i.e., $C = I_{n_x}$,*

then, for $K \in \mathcal{K}$, there exists a constant $\mu_K > 0$ such that

$$\mu_K (J(K) - J^*) \leq \text{dist}(0, \partial J(K)), \quad (16)$$

where ∂J denotes the Clarke subdifferential of J and $J^* = \min_{K \in \mathcal{K}} J(K)$.

The proof is provided in [Section B](#). We also provide an explicit form of the constant μ_K in (16) in the proof. This remarkable property has a tight connection to the well-known LMI reformulation via the bounded real lemma [30], which will be utilized in the proof. Our proof is motivated by the recent ECL framework [6] and the result in [24]. The global optimality of stationary points follows as a corollary.

Corollary 1 (Global optimality of stationary points). *Under the same assumptions as [Theorem 2](#), we have $0 \in \partial J(K) \Leftrightarrow K \in \arg \min_{K \in \mathcal{K}} J(K)$.*

Proof. If $0 \in \partial J(K)$, then we have $\text{dist}(0, \partial J(K)) = 0$, which implies that $J(K) = J^*$ by the weak PL (16). On the other hand, if $K \in \arg \min_{K \in \mathcal{K}} J(K)$, we must have K to be stationary, and thus $0 \in \partial J(K)$ [23, Th. 10.1]. $\square$

D. Proof of [Theorem 1](#)

To establish [Theorem 1](#), we first rewrite the $\mathcal{H}_\infty$ cost as the composition of $\sigma_{\max}(\cdot)$ and a real-valued mapping using [Lemma 4](#) below, which follows from a famous fact [29, Exercise 4.42] in complex SDPs. We thereby identify the weakly convex constant $m_\nu > 0$ via the compactness of $\mathcal{K}_\nu$ from [Lemma 1](#). The proof of [Lemma 4](#) is presented in [Section A](#) due to the page limit.

Lemma 4. *Let X be a $q \times r$ hermitian matrix. We have*

$$\sigma^2 I \succeq X^H X \iff \sigma I \succeq \Psi(X)$$

where $\Psi(\cdot)$ is a real-valued linear mapping defined by

$$\Psi(X) = \begin{bmatrix} \begin{bmatrix} 0 & \Re(X)^T \\ \Re(X) & 0 \end{bmatrix} & \begin{bmatrix} 0 & \Im(X)^T \\ -\Im(X) & 0 \end{bmatrix} \\ \begin{bmatrix} 0 & -\Im(X)^T \\ \Im(X) & 0 \end{bmatrix} & \begin{bmatrix} 0 & \Re(X)^T \\ \Re(X) & 0 \end{bmatrix} \end{bmatrix},$$

where $\Re(X)$ and $\Im(X)$ represent the real and imaginary parts, respectively.

Recall the definition of the function $J(K) = \|\mathbf{T}_{zw}(K)\|_{\mathcal{H}_\infty}$ in (7). We first show the weak convexity of

$$f(\cdot, \omega) := \sigma_{\max}(\mathbf{T}_{zw}(\cdot, \omega))$$

for any fixed $\omega \in [0, 2\pi]$. Then, the weak convexity of $J = \max_{\omega \in [0, 2\pi]} f(\cdot, \omega)$ follows from the standard result for point-wise maximum of convex functions [29, Sec. 3.2.3].

Fix an $\omega \in [0, 2\pi]$. By applying [Lemma 4](#), we have that

$$f(K, \omega) = \sigma_{\max}(\mathbf{T}_{zw}(K, \omega)) = \lambda_{\max}((\Psi \circ \mathbf{T}_{zw})(K, \omega)).$$

It is clear that both $\Re(\mathbf{T}_{zw}(K, \omega))$ and $\Im(\mathbf{T}_{zw}(K, \omega))$ are real analytic on $\mathcal{K} \times [0, 2\pi]$. Thus, the mapping $\Psi \circ \mathbf{T}_{zw}$ is also real analytical over $\mathcal{K} \times [0, 2\pi]$. Then, the Hessian of $\Psi \circ \mathbf{T}_{zw}$ exists and is continuous w.r.t. (K, ω) over $\mathcal{K} \times [0, 2\pi]$.

By the compactness of $\mathcal{K}_\nu$, we have the L_ν -smoothness of $f(\cdot, \omega)$ on $\mathcal{K}_\nu$ with the following finite positive constant:

$$L_\nu = 4\eta^2 \max_{(i,j)} \max_{\substack{K \in \mathcal{K}_\nu, \\ \omega \in [0, 2\pi]}} \left\| \nabla_{\text{vec}(K)}^2 [(\Psi \circ \mathbf{T}_{zw})]_{ij}(K, \omega) \right\|_F,$$

where $\eta = n_x + n_u + n_w$. By recalling that $\lambda_{\max}(\cdot)$ is convex and 1-Lipschitz, $f(K, \omega)$ is m_ν -weakly convex over any convex subset of $\mathcal{K}_\nu$ with $m_\nu = L_\nu$; see [18, Sec. 2.1].

Finally, it is straightforward to see that the function

$$J(K) + \frac{m_\nu}{2} \|K\|_F^2 = \max_{\omega \in [0, 2\pi]} \left\{ f(K, \omega) + \frac{m_\nu}{2} \|K\|_F^2 \right\}$$

is convex on any convex subset of $\mathcal{K}_\nu$ from the m_ν -weak convexity of each $f(K, \omega)$ [29, Sec. 3.2.3]. This leads to the desired m_ν -weak convexity of $J(K)$. $\square$

IV. THE SUBGRADIENT METHOD

In this section, we discuss the subgradient method to solve the $\mathcal{H}_\infty$ optimization (8). Thanks to the weak convexity in [Theorem 1](#), we provide a non-asymptotic convergence rate under mild assumptions, by adapting the recent result in [18].

A. The subgradient method

Given an initial point $K_0 \in \mathcal{K}$, the subgradient method simply follows the update below

$$K_{t+1} = K_t - \alpha_t G_t, \quad G_t \in \partial J(K_t), \quad (17)$$

where $\partial J(K)$ is the Clarke subdifferential of $J(\cdot)$ at $K \in \mathcal{K}$ [23]. This is arguably the simplest algorithm for nonsmooth nonconvex optimization. Thanks to its simplicity, we may implement (17) in a model-free manner relatively easily, only with a subgradient oracle.

B. Non-asymptotic convergence guarantees

Despite its simplicity, the subgradient method is classically known to achieve sublinear convergence only for convex and/or smooth problems [16], [18], [31]. Its behavior in the nonsmooth nonconvex setting was far less understood because (i) stationary concepts in convex and/or smooth cases can no longer be used [16], [18], and (ii) the subgradient update (17) does not guarantee monotonic decrease of the cost; it is indeed possible even in the nonsmooth convex case that $J(K_{t+1}) > J(K_t)$ for any $\alpha_t > 0$. Recently, [18] established the first non-asymptotic rate for minimizing *unconstrained* weakly convex functions. In this work, we extend their analysis to the $\mathcal{H}_\infty$ optimization (8) involving a *nonconvex stabilizing constraint* $\mathcal{K}$.

One key concept is the *Moreau envelope*, which allows for quantifying the performance of (17) even in the absence of a function value decrease.

Definition 2 (Moreau envelope). *For $K \in \mathcal{K}$ and $\rho > 0$, the Moreau envelope $J_\rho(\cdot)$ is defined as $J_\rho(K) := \min_{M \in \mathcal{K}} J(M) + \frac{\rho}{2} \|M - K\|_F^2$.*

By the coercivity in [Lemma 1](#), there always exists a minimizer $\hat{K} \in \arg \min_{M \in \mathcal{K}} J(M) + \frac{\rho}{2} \|M - K\|_F^2$, but the minimizer may not be unique. By the definition, we have

$$J_\rho(K) \leq J(K), \text{ and } J_\rho(\hat{K}) = J(\hat{K}) \leq J(K), \quad \forall K \in \mathcal{K}.$$

While the Moreau envelope is convenient in theory, computing the value of $J_\rho(\cdot)$ and $\hat{K}$ might be untractable in practice. If $\rho > 0$ is sufficiently large, the minimizer $\hat{K}$ will be unique. We have the following result by adapting [18, Sec. 2.2].

Lemma 5. Consider the $\mathcal{H}_\infty$ cost function $J : \mathcal{K} \rightarrow \mathbb{R}_+$ defined in (8). Fix any $K \in \mathcal{K}$. Let $\nu > J(K)$, and define $L > 0$ and $m > 0$ such that $J(\cdot)$ is L -Lipschitz over $\mathcal{K}_\nu$ and m -weakly convex over any convex subset of $\mathcal{K}_\nu$. Then, for $\rho > \max\{m, L/\text{dist}(K, \partial\mathcal{K}_\nu)\}$, we have

$$\begin{aligned} \|K - \hat{K}\|_F &\leq \text{dist}(K, \partial\mathcal{K}_\nu) \\ \nabla J_\rho(K) &= \rho(K - \hat{K}), \\ \text{dist}(0, \partial J(\hat{K})) &\leq \|\nabla J_\rho(K)\|_F, \end{aligned}$$

where $\{\hat{K}\} = \arg \min_{M \in \mathcal{K}} J(M) + \frac{\rho}{2} \|M - K\|_F^2$.

Proof. From the optimality condition of $\min_M J(M) + \frac{\rho}{2} \|M - K\|_F^2$, we have

$$\rho(K - \hat{K}) \in \partial J(\hat{K}), \quad (18)$$

where $\hat{K}$ is any minimizer. The L -Lipschitzness of J implies $\|K - \hat{K}\|_F \leq L/\rho \leq d := \text{dist}(K, \partial\mathcal{K}_\nu)$, where the last inequality is by the choice of $\rho > \max\{m, L/\text{dist}(K, \partial\mathcal{K}_\nu)\}$. Thus, minimizing J over $\mathcal{K}$ is the same as minimizing J over the ball $\mathbb{B}_d(K)$, i.e., $\arg \min_{M \in \mathcal{K}} \{J(M) + \frac{\rho}{2} \|M - K\|_F^2\} = \arg \min_{M \in \mathbb{B}_d(K)} \{J(M) + \frac{\rho}{2} \|M - K\|_F^2\}$.

Since $\mathbb{B}_d(K) \subseteq \mathcal{K}_\nu$ and $\rho > m$, we have $J(\cdot) + \frac{\rho}{2} \|\cdot - K\|_F^2$ being $(\rho - m)$ -strongly convex over $\mathbb{B}_d(K)$. The minimizer is then unique. Hence, from [18, Lem. 2.2], the gradient of J_ρ can be computed as $\nabla J_\rho(K) = \rho(K - \hat{K})$, and (18) shows $\text{dist}(0, \partial J(\hat{K})) \leq \|\nabla J_\rho(K)\|_F$. $\square$

We are now ready to state the convergence of (17) in terms of the Moreau envelope under a technical condition that the costs are bounded along the iterations.

Theorem 3. Consider the $\mathcal{H}_\infty$ cost function $J : \mathcal{K} \rightarrow \mathbb{R}_+$ defined in (8). Let $\{K_0, \dots, K_{T+1}\}$ be a sequence generated by the subgradient method (17) with $K_0 \in \mathcal{K}$ and step sizes $\{\alpha_0, \dots, \alpha_T\}$. Suppose $\max_{t \in \{0, \dots, T+1\}} J(K_t) < \bar{J} < \infty$ and $\max_{t \in \{0, \dots, T\}} \|G_t\|_F \leq H < \infty$. Then, we have the feasibility of the iterates $\{K_t\} \subset \mathcal{K}$, and

$$\min_{t=0, \dots, T} \|\nabla J_\rho(K_t)\|_F^2 \leq \kappa \frac{J_\rho(K_0) - J^* + \frac{\rho_{\min} H^2}{2} \sum_{t=0}^T \alpha_t^2}{\sum_{t=0}^T \alpha_t},$$


where $\kappa = \rho/(\rho - m)$ with some constant $m > 0$, and $\rho > \rho_{\min} \in (m, \infty)$. Furthermore, if $\alpha_0 = \dots = \alpha_T = \alpha = \beta/\sqrt{T+1}$ with $\beta > 0$, we have

$$\min_{t=0, \dots, T} \|\nabla J_\rho(K_t)\|_F^2 \leq \kappa \frac{J_\rho(K_0) - J^* + \rho_{\min} H^2 \beta^2 / 2}{\beta \sqrt{T+1}}.$$

Proof. Consider the sublevel set $\mathcal{K}_{\bar{J}} = \{K \in \mathcal{K} \mid J(K) \leq \bar{J}\}$, which is compact due to the coercivity. By Theorem 1 and Lemma 1, $J(\cdot)$ is L -Lipschitz over $\mathcal{K}_{\bar{J}}$ and m -weakly convex over any convex subset of $\mathcal{K}_{\bar{J}}$, with some $L, m > 0$.

Let $\rho_{\min} > \max\{m, L/\tilde{d}\}$, where

$$\tilde{d} := \min_{t \in \{0, \dots, T+1\}} \text{dist}(K_t, \partial\mathcal{K}_{\bar{J}}) > 0.$$



(a) System (10) ($C = I_{n_x}$)

(b) Example 4 ($C \neq I_{n_x}$)

Fig. 4: The simulation results of the subgradient method (17) with $\alpha_t = \alpha^1 = 10^{-3}$ or $\alpha^2 = 10^{-4}$: (a) The plot of $\min_t |J(K_t) - J^*|/J^*$ for the system (10) with a full state-feedback in Example 1 starting from $K_0 = K_0^1$ or K_0^2 , where $K_0^1 = [0, -1.9]$ and $K_0^2 = [0.2, -2]$; (b) The plots of $\min_t J(K_t)$ (upper) and $\min_t \|G_t\|_F^2$ (lower) for Example 4 with $K_0 = K_0^1 = [0, 0]$ or $K_0^2 = [-5, -2]$.

Then, the same argument as Lemma 5 implies that $\mathbb{B}_{\tilde{d}}(K_t) \subset \mathcal{K}_{\bar{J}}$ for all t , and $\hat{K}_t \in \arg \min_M J(M) + \frac{\rho}{2} \|M - K_t\|_F^2$ is contained by the ball $\mathbb{B}_{\tilde{d}}(K_t)$. Hence, we can regard $J(K) + \frac{\rho}{2} \|K - K_t\|_F^2$ as a $(\rho - m)$ -strongly convex function over $\mathbb{B}_{\tilde{d}}(K_t)$, and this theorem follows from performing the same argument as [18, Th. 3.1]. The coercivity in Lemma 1 yields the feasibility of $\{K_t\}$. $\square$

This theorem indicates that we can find an iterate K_{t^*} satisfying $\|\nabla J_\rho(K_{t^*})\| \leq \epsilon$ after

$$T = \left\lceil \kappa^2 \left(J_\rho(K_0) - J^* + \frac{m H^2 \beta^2}{2} \right)^2 \beta^{-2} \epsilon^{-4} - 1 \right\rceil = \mathcal{O}(\epsilon^{-4})$$

iterations, as long as the costs remain finite. For K_{t^*} ,

$$\text{dist}(0, \partial J(\hat{K}_{t^*})) \leq \epsilon, \quad \|\hat{K}_{t^*} - K_{t^*}\|_F \leq \epsilon/\rho$$

hold, where $\hat{K}_{t^*} = \arg \min_{M \in \mathcal{K}} J(M) + \frac{\rho}{2} \|M - K_{t^*}\|_F^2$. Thus, the subgradient method (17) allows us to arrive at an approximation of a stationary point in the above sense.

Remark 2. Theorem 3 presents a non-asymptotic convergence rate for the subgradient method (17). Notice that the boundedness of the cost is directly assumed, as it is often empirically true for sufficiently small $\{\alpha_t\}$ but not formally guaranteed in (17). Under this assumption, we can ensure the feasibility $\{K_t\} \subset \mathcal{K}$ by the coercivity in Lemma 1, bypassing the nonconvex constraint set $\mathcal{K}$ through a careful choice of the parameter ρ . To alleviate the assumption on the cost, one may use other algorithms that can generate a descent direction (e.g., the proximal bundle method [21]). We leave it as our future work. $\square$

V. NUMERICAL RESULTS

Here, we run the subgradient method (17) with fixed step sizes $\alpha_t = \alpha^1$ or α^2 , where $(\alpha^1, \alpha^2) = (10^{-3}, 10^{-4})$, for both full state-feedback and static output-feedback cases. The following results demonstrate the effectiveness of the subgradient method (17) and validate our theoretical results.

Full state-feedback. For the system (10) in Example 1, we test (17) with different parameter choices. Specifically, we use $K_0 = K_0^1 = [0, -1.9]$ or $K_0^2 = [0.2, -2]$. We plot the simulation result in Figure 4a, where the vertical axis shows the best relative objective residual $\min_t |J(K_t) - J^*|/J^*$ at t . In all the cases, the subgradient method generates feasible iterates, which implies that the assumptions in Theorem 3 are satisfied. In addition, as expected, we obtained a globally optimal solution for both initial conditions with favorable accuracy, especially when $\alpha_t = \alpha^2$.

Static output-feedback. For Example 4, we similarly run (17) with four different choices of the initial condition and step size. In particular, we use $K_0 = K_0^1$ or K_0^2 with $K_0^1 = [0, 0]$, $K_0^2 = [-5, -2]$. The simulation result is presented in Figure 4b. It can be seen that the feasibility is satisfied in all the cases. Regarding the performance, the cost value converged to a smaller value in the case of $K_0 = K_0^1$. In addition, when $K_0 = K_0^2$, we observe that the subgradient also converged nearly to a stationary point. When $K_0 = K_0^1$ in Figure 4b, the magnitude of the subgradient remained larger despite the smaller cost value. This is due to a zigzagging behavior around a stationary point and may happen when it is nonsmooth. For example, $f(x) = |x|$ takes the minimum at $x = 0$, but the origin is nonsmooth and $\partial f(0) = [-1, 1] \neq \{0\}$, which may lead to such behavior.

VI. CONCLUSION

This study addressed the $\mathcal{H}_\infty$ optimization problem for discrete-time linear systems with static output-feedback. We first proved the weak convexity and further established the weak PL condition in the full state-feedback case. Then, we also discussed the subgradient method from a weak convexity perspective, presenting the first non-asymptotic convergence rate. The numerical simulations validated our theoretical analysis and demonstrated the effectiveness of the subgradient method. Our future work includes alleviating the assumption of the cost value boundedness in the convergence rate result, for example, through the proximal bundle method.

REFERENCES

- [1] K. Zhou, J. C. Doyle, and K. Glover, *Robust and Optimal Control*. Prentice Hall, 1996.
- [2] P. Apkarian and D. Noll, "Nonsmooth $\mathcal{H}_\infty$ synthesis," *IEEE Transactions on Automatic Control*, vol. 51, no. 1, pp. 71–86, 2006.
- [3] X. Guo and B. Hu, "Global convergence of direct policy search for state-feedback $\mathcal{H}_\infty$ robust control: A revisit of nonsmooth synthesis with goldstein subdifferential," *Advances in Neural Information Processing Systems*, vol. 35, pp. 32 801–32 815, 2022.
- [4] X. Guo, D. Keivan, G. Dullerud, P. Seiler, and B. Hu, "Complexity of derivative-free policy optimization for structured $\mathcal{H}_\infty$ control," *Advances in Neural Information Processing Systems*, vol. 36, pp. 5050–5078, 2023.
- [5] Y. Tang and Y. Zheng, "On the global optimality of direct policy search for nonsmooth $\mathcal{H}_\infty$ output-feedback control," in *2023 62nd IEEE Conference on Decision and Control (CDC)*. IEEE, 2023, pp. 6148–6153.
- [6] Y. Zheng, C.-F. Pai, and Y. Tang, "Benign nonconvex landscapes in optimal and robust control, Part I: Global optimality," *arXiv preprint arXiv:2312.15332*, 2023.
- [7] Y. Zheng, C.-F. Pai, and Y. Tang, "Benign nonconvex landscapes in optimal and robust control, Part II: Extended Convex Lifting," *arXiv preprint arXiv:2406.04001*, 2024.

- [8] B. Hu and Y. Zheng, "Connectivity of the feasible and sublevel sets of dynamic output feedback control with robustness constraints," *IEEE Control Systems Letters*, vol. 7, pp. 442–447, 2022.
- [9] J. C. Doyle, K. Glover, P. Khargonekar, and B. Francis, "State-space solutions to standard H_2 and H_∞ control problems," in *1988 American control conference*. IEEE, 1988, pp. 1691–1696.
- [10] P. Gahinet and P. Apkarian, "A linear matrix inequality approach to H_∞ control," *International Journal of Robust and Nonlinear Control*, vol. 4, no. 4, pp. 421–448, 1994.
- [11] J. V. Burke, D. Henrion, A. S. Lewis, and M. L. Overton, "HIFOO—a MATLAB package for fixed-order controller design and $\mathcal{H}_\infty$ optimization," *IFAC Proceedings Volumes*, vol. 39, no. 9, pp. 339–344, 2006.
- [12] M. Saeki, "Fixed structure PID controller design for standard H_∞ control problem," *Automatica*, vol. 42, no. 1, pp. 93–100, 2006.
- [13] P. Apkarian, D. Noll, and O. Prot, "A proximity control algorithm to minimize nonsmooth and nonconvex semi-infinite maximum eigenvalue functions," *Journal of Convex Analysis*, vol. 16, no. 3–4, pp. 641–666, 2009.
- [14] P. Apkarian and D. Noll, "The H_∞ control problem is solved," *Aerospace Lab*, no. 13, pp. pages–1, 2017.
- [15] J. V. Burke, A. S. Lewis, and M. L. Overton, "A robust gradient sampling algorithm for nonsmooth, nonconvex optimization," *SIAM Journal on Optimization*, vol. 15, no. 3, pp. 751–779, 2005.
- [16] J. Zhang, H. Lin, S. Jegelka, S. Sra, and A. Jadbabaie, "Complexity of finding stationary points of nonconvex nonsmooth functions," in *International Conference on Machine Learning*. PMLR, 2020, pp. 11 173–11 182.
- [17] J. C. Duchi, P. L. Bartlett, and M. J. Wainwright, "Randomized smoothing for stochastic optimization," *SIAM Journal on Optimization*, vol. 22, no. 2, pp. 674–701, 2012.
- [18] D. Davis and D. Drusvyatskiy, "Stochastic model-based minimization of weakly convex functions," *SIAM Journal on Optimization*, vol. 29, no. 1, pp. 207–239, 2019.
- [19] J. Liang, R. D. Monteiro, and H. Zhang, "Proximal bundle methods for hybrid weakly convex composite optimization problems," *arXiv preprint arXiv:2303.14896*, 2023.
- [20] F.-Y. Liao, L. Ding, and Y. Zheng, "Error bounds, PL condition, and quadratic growth for weakly convex functions, and linear convergences of proximal point methods," in *6th Annual Learning for Dynamics & Control Conference*. PMLR, 2024, pp. 993–1005.
- [21] F.-Y. Liao and Y. Zheng, "A proximal descent method for minimizing weakly convex optimization," *arXiv preprint arXiv:2509.02804*, 2025.
- [22] D. Noll, "Bundle method for non-convex minimization with inexact subgradients and function values," in *Computational and Analytical Mathematics: In Honor of Jonathan Borwein's 60th Birthday*. Springer, 2013, pp. 555–592.
- [23] R. T. Rockafellar and R. J.-B. Wets, *Variational analysis*. Springer Science & Business Media, 2009, vol. 317.
- [24] J. Umenberger, M. Simchowitz, J. Perdomo, K. Zhang, and R. Tedrake, "Globally convergent policy search for output estimation," *Advances in Neural Information Processing Systems*, vol. 35, pp. 22 778–22 790, 2022.
- [25] K. Glover and J. C. Doyle, "A state space approach to $\mathcal{H}_\infty$ optimal control," in *Three Decades of Mathematical System Theory: A Collection of Surveys at the Occasion of the 50th Birthday of Jan C. Willems*, H. Nijmeijer and J. M. Schumacher, Eds. Springer-Verlag, 1989, pp. 179–218.
- [26] B. Hu, K. Zhang, N. Li, M. Mesbahi, M. Fazel, and T. Başar, "Toward a theoretical foundation of policy optimization for learning control policies," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 6, no. 1, pp. 123–158, 2023.
- [27] S. Talebi, Y. Zheng, S. Kraissler, N. Li, and M. Mesbahi, "Policy optimization in control: Geometry and algorithmic implications," *arXiv preprint arXiv:2406.04243*, 2024.
- [28] H. Feng and J. Lavaei, "Connectivity properties of the set of stabilizing static decentralized controllers," *SIAM Journal on Control and Optimization*, vol. 58, no. 5, pp. 2790–2820, 2020.
- [29] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [30] S. P. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*. SIAM, 1994.
- [31] Y. Nesterov, *Lectures on convex optimization*. Springer, 2018, vol. 137.

A. Proof of Lemma 4

Consider $X \neq 0$ without loss of generality. By the Schur complement, we have

$$\sigma I \succeq \frac{1}{\sigma} X^H X \iff \sigma I - \begin{bmatrix} 0 & X^H \\ X & 0 \end{bmatrix} \succeq 0. \quad (19)$$

Now, from [29, Exercise 4.42] for a hermitian matrix Z , we know that $\lambda I \succeq Z$ is equivalent to $\lambda I \succeq \begin{bmatrix} \Re(Z) & -\Im(Z) \\ \Im(Z) & \Re(Z) \end{bmatrix}$.

With this, letting $Y = \begin{bmatrix} 0 & X^H \\ X & 0 \end{bmatrix}$ in (19), we obtain

$$\sigma I - Y \succeq 0 \iff \sigma I \succeq \begin{bmatrix} \Re(Y) & -\Im(Y) \\ \Im(Y) & \Re(Y) \end{bmatrix} = \Psi(X).$$

This completes the proof. $\square$

B. Proof of Theorem 2

We derive the weak PL condition (16) by exploiting a favorable *partial minimization* structure of the cost J . First, by using the nonstrict bounded real lemma [30, Sec. 2.3.7], it is known that the problem (8) with $C = I_{n_x}$ admits the following reformulation with the additional Lyapunov variable P and a scalar γ :

$$\min_{\gamma, K, X} \gamma \quad \text{subject to} \quad (K, \gamma, P) \in \mathcal{L}_{\text{lift}}, \quad (20)$$

where we define the set $\mathcal{L}_{\text{lift}}$ as

$$\mathcal{L}_{\text{lift}} = \{(K, \gamma, P) \mid P \succ 0, \Lambda(K, \gamma, P) \preceq 0\}$$

with

$$\Lambda(K, \gamma, P) = \begin{bmatrix} A_K^T P A_K - P & A_K^T P \\ P A_K & P \end{bmatrix} + \begin{bmatrix} Q + K^T R K & 0 \\ 0 & -\gamma^2 I \end{bmatrix}$$

where $A_K = A + BK$. Then, we can represent $J(K)$ as the partial minimization of the following lifted function:

$$J_{\text{lift}}(K, \gamma, P) = \gamma + \delta_{\mathcal{L}_{\text{lift}}}(K, \gamma, P)$$

with the indicator function $\delta_{\mathcal{L}_{\text{lift}}}(K, \gamma, P)$ for $\mathcal{L}_{\text{lift}}$. Namely,

$$J(K) = \inf_{\gamma, P} J_{\text{lift}}(K, \gamma, P). \quad (21)$$

We first show the weak PL inequality for $J_{\text{lift}}(K, \gamma, P)$, by which we derive (16). We here use the exact convex reformulation (LMI) [30] for (20) with the change of variables $X = (P/\gamma)^{-1}$ and $Y = KX$. To handle this, we define the following C^∞ mapping $\Pi : \mathcal{L}_{\text{lift}} \rightarrow \mathcal{F}_{\text{cvx}}$:

$$\Pi(K, \gamma, P) = (\gamma, K(P/\gamma)^{-1}, (P/\gamma)^{-1}),$$

with the following convex set

$$\mathcal{F}_{\text{cvx}} = \{(\gamma, Y, X) \mid X \succ 0, \text{LMI}(\gamma, Y, X) \preceq 0\},$$

where $\text{LMI}(\gamma, Y, X)$ is the following linear map:

$$\text{LMI}(\gamma, Y, X) = \begin{bmatrix} -X & 0 & X & (AX + BY)^T & Y^T \\ 0 & -\gamma I & 0 & I & 0 \\ X & 0 & -\gamma Q^{-1} & 0 & 0 \\ (AX + BY) & I & 0 & -X & 0 \\ Y & 0 & 0 & 0 & -\gamma R^{-1} \end{bmatrix}.$$

This LMI is standard and follows from the inequalities $\Lambda(K, \gamma, P) \preceq 0$, $P \succ 0$ and Schur complement, which yields

$$\Pi(\mathcal{L}_{\text{lift}}) = \mathcal{F}_{\text{cvx}}.$$

Notice that this mapping $\Pi(\cdot, \cdot, \cdot)$ is a diffeomorphism. Then, by defining the function

$$J_{\text{cvx}}(\gamma, Y, X) = \gamma + \delta_{\mathcal{F}_{\text{cvx}}}(\gamma, Y, X),$$

which is convex from the convexity of $\mathcal{F}_{\text{cvx}}$, we can further rewrite

$$J_{\text{lift}}(K, \gamma, P) = (J_{\text{cvx}} \circ \Pi)(K, \gamma, P).$$

Note also that $J_{\text{cvx}}(\gamma, Y, X) = (J_{\text{lift}} \circ \Pi^{-1})(\gamma, Y, X)$.

Now we are ready to establish the weak PL condition. By the convexity of J_{cvx} ,

$$\begin{aligned} J_{\text{lift}}(K, \gamma, P) - J^* &= (J_{\text{cvx}} \circ \Pi)(K, \gamma, P) - J^* \\ &\leq \langle W, \Pi(K, \gamma, P) - \Pi(K^*, J^*, P^*) \rangle \\ &\leq \|W\|_F \times \|\Pi(K, \gamma, P) - \Pi(K^*, J^*, P^*)\|_F \end{aligned} \quad (22)$$

where $W \in \partial J_{\text{cvx}}(\Pi(K, \gamma, P))$, and (K^*, J^*, P^*) represents an optimal triplet. The chain rule [23, Th. 10.6] for $J_{\text{cvx}} = J_{\text{lift}} \circ \Pi^{-1}$ yields

$$J_{\text{lift}}(K, \gamma, P) - J^* \leq \|V\|_F \times \beta_{(K, \gamma, P)}, \quad (23)$$

where $V \in \partial J_{\text{lift}}(K, \gamma, P)$ and

$$\begin{aligned} \beta_{(K, \gamma, P)} &= \sigma_{\max}(\nabla \Pi^{-1}(\Pi(K, \gamma, P))) \\ &\quad \times \|\Pi(K, \gamma, P) - \Pi(K^*, J^*, P^*)\|_F > 0. \end{aligned}$$

Note that $\beta_{(K, \gamma, P)}$ is positive for $(K, \gamma, P) \in \mathcal{L}_{\text{lift}}$. Recalling the partial minimization property of $J(K)$ in (21), we now have the following fact from [23, Th. 10.13]:

$$\partial J(K) \subset \bigcup_{(\gamma, P) \in \arg \inf_{\gamma, P} J_{\text{lift}}(K, \gamma, P)} \mathcal{M}(\gamma, P),$$

where

$$\mathcal{M}(\gamma, P) = \{G \mid (G, 0, 0) \in \partial J_{\text{lift}}(K, \gamma, P)\}.$$

Note that the assumptions in [23, Th. 10.13] are satisfied due to the boundedness of the level set; whenever $\|K\|_F \rightarrow \infty$ or $\|P\|_F \rightarrow \infty$, it leads to the infeasibility, and $J_{\text{lift}}(K, \gamma, P) \rightarrow \infty$, which implies the level set $\{(K, \gamma, P) \in \mathcal{L}_{\text{lift}} \mid J_{\text{lift}}(K, \gamma, P) \leq t\}$ for any $t \geq 0$ is bounded.

Then substituting $(\gamma_K, P_K) \in \arg \inf_{\gamma, P} J_{\text{lift}}(K, \gamma, P)$ into (23) gives

$$J(K) - J^* \leq \|G\|_F \times \beta_{(K, \gamma_K, P_K)},$$

where G can be taken as any subgradient in $\partial J(K)$ thanks to the arbitrariness of $W \in \partial_{(\gamma, Y, X)} J_{\text{cvx}}(\Pi(K, \gamma, P))$ in (22). Thus, we arrive at the weak PL inequality in (16) with

$$\mu_K = \max_{(\gamma_K, P_K) \in \arg \min_{\gamma, P} J_{\text{lift}}(K, \gamma, P)} \beta_{(K, \gamma_K, P_K)}^{-1} > 0, \quad (24)$$

where μ_K is positive as $K \in \mathcal{K}$. $\square$